

PROTOCOLS FOR ANONYMOUS ELECTRONIC COMMUNICATION AND DOUBLE-BLIND TRANSACTIONS

DESCRIPTION

5

1

BACKGROUND OF THE INVENTION

Field of the Invention

09550462-041700

10 The present invention generally relates to electronic commerce on a distributed computer network, such as the Internet, and more particularly, to a system and associated protocols for communication between two entities across a computer network such that their identities remain concealed from each other, while ensuring that no third party is able to trace the existence of a conversation between them. The invention aims to ensure that the anonymity thus provided does not depend on the existence of any single trusted agent (or

15 trusted third party) and is not compromised by the existence of malicious agents within or outside the system.

Background Description

20 In U.S. Patent No. 5,794,207 to Walker et al., there is proposed a method and an electronic apparatus that allows prospective buyers of goods and services to communicate a binding purchase offer globally to potential sellers, for sellers to conveniently search for relevant purchase offers, and for sellers potentially to bind a buyer to a contract based on the buyer's purchase

However, in this emerging world of e-commerce, the issue of privacy is believed by many to be a significant impediment to future growth.

David Chaum in an article entitled “Untraceable electronic mail, return addresses, and digital pseudonyms”, *Communications of the ACM*, 24, 2 (February 1981), pp. 84–88, proposed a system for anonymous electronic mail which employs a set of forwarding agents called mixes. The routing of messages through the various mixes was performed using a source routing algorithm. Each mix collects a few messages, waits a period of time and sends the messages out in a different order. Mixes are meant to prevent global eavesdroppers from tracing messages passing through them and thus provide sender and receiver unlinkability. But since routing is done at the source, there is no receiver anonymity. The source knows the receiver’s identity. The strength of mixes is that even if one mix in a path is not compromised, the system continues to provide sender-receiver unlinkability against global eavesdroppers.

M. K. Reiter and A. D. Rubin in an article entitled “Crowds: Anonymity for Web Transactions”, *ACM Transactions on Information and System Security*, describes another approach. Crowds is a system consisting of a cooperative group of users which provides sender anonymity. It also provides receiver anonymity against local eavesdroppers and colluding forwarders called “jondos”. But Crowds does not envisage providing receiver anonymity against the sender himself. The Crowds system is basically designed for anonymous access to web pages where it is understood that the sender knows the location of the page he wants to access.

SUMMARY OF THE INVENTION

It is therefore an object of the invention to provide a system and associated protocols for communication between two entities across a computer network such that their identities remain concealed from each other, while ensuring that no third party is able to trace the existence of a conversation between them.

nsa
0950462-04-1700

The protocols of the present invention for communication between two entities across a network are such that their identities remain concealed from each other, while ensuring that no third party is able to trace the existence of a conversation between them. The two entities correspond to each other through pseudonyms. The protocols are also designed with an object to distribute trust so that an identity is not revealed by the compromise of any one agent involved in the execution of the protocol. No one agent can establish a correlation between a pseudonym and a physical address.

In contrast to mixes as described by Chaum, *supra*, our system is meant to provide receiver anonymity and does not attempt to defend itself against a global eavesdropper by direct means. The probability of compromise of identity in our system is a polynomial function of the number of compromised agents. Therefore, our system promises that the expected number of compromises will not be proportional to the number of compromised agents. The loss would be less than proportional when the number of colluders is low.

In contrast to crowds described by Reiter and Rubin, *supra*, our system provides receiver anonymity against the sender himself by making him communicate with a third party and addressing the message to a pseudonym.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

5 Figure 1 is a flow diagram of the process of registration of an identity in the system according to the present invention;

 Figure 2 is a flow diagram showing the implementation of the forwarding algorithm according to the present invention;

10 Figure 3 is a flow diagram of the setup process for modification of a message at the originating agent's end as performed in the process of Figure 2;

 Figures 4A and 4B, taken together, are a flow diagram of the consistency verification as performed in the process of Figure 2 according to the solution one protocol of the present invention;

15 Figure 5 is a flow diagram of the process of choosing an agent and modifying a header as performed in the process of Figure 2 according to the solution one protocol of the present invention;

 Figure 6 is a flow diagram of the process of consistency verification in the process of Figure 2 according to the solution two protocol of the present invention;

20 Figures 7A and 7B, taken together, are a flow diagram of the process of consistency verification in the solution two protocol executed by coordinating agents (CAs) of the present invention;

 Figure 8 is a flow diagram of protocol consistency in the solution two protocol of the present invention; and

25 Figure 9 is a flow diagram of the process of choosing an agent and modifying a header in the solution two protocol according to the invention.

00550462-04-1700

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS OF THE INVENTION

There are essentially two aspects to double blind transactions, sender anonymity and receiver anonymity. There have been many approaches proposed for sender anonymity, among which are mixes (see D. Chaum) and Crowds (see Reiter and Rubin). Receiver anonymity is provided by a technique described in detail in the following section. Essentially, there are a set of forwarding agents who, individually, do not possess sufficient information to compromise an identity, but can collectively forward a message to the right destination.

The System

The system consists of two parts, a set of clients and a set of Forwarding Agents (FAs). Let there be n FAs, say F_1, F_2, \dots, F_n . Several groups of these n agents, each of which consists of k members, are listed, where k ($0 < k \leq n$) is a fixed number considered sufficient to provide anonymity in the system. Each FA must belong to at least one group, but there is no restriction on the number of groups that an FA can belong to. Each of the FAs has its own pair of public and private keys for encryption and decryption, respectively, where the underlying cryptosystem scheme is some *commutative* public key cryptosystem (a *commutative cryptosystem* is a cryptosystem in which multiple decryptions, possibly with different keys, performed on a ciphertext produce the same result irrespective of the order in which such decryptions are done). One such cryptosystem is the RSA (Rivest, Shamir and Adleman) cryptosystem. Each FA also has the appropriate keys required to perform secure digital signatures on documents and to verify the signatures of

other FAs.

As shown in Figure 1, each client of the system has to initially register with a Forwarding Agent S of his or her choice in function block 100. The client, having selected a Forwarding Agent S, also picks one of the groups that the Forwarding Agent S belongs to, thus selecting k agents to be associated with the client. This registration process, as described in the next section, involves the assignment of a pseudonym to the client in function block 150. The client also provides the Forwarding Agent S with an encrypted form of his or her network address, rendering it unreadable to any individual FA, as described in the next section. Each FA maintains a table with three fields, namely a pseudonym, a corresponding encrypted network address and the FA group to be used for forwarding.

Informally, the system delivers a message as follows. A message meant for a pseudonym X arrives at the FA where X is registered. The message then goes through a random sequence of FAs. The set of FAs through which the message passes is fixed for each pseudonym. The last FA in the sequence finds a visible network address and sends the message on to this address. The next section describes this process in detail.

We describe below two variations of our protocol, which we refer to as Solution One and Solution Two, respectively. These two protocols are largely identical but differ in some steps and computations performed by agents. Below, we first describe the basic structure of the protocol which is common to both variations. The details in which they differ are then described separately for both.

5

10

15

20

25

the system and has publicized its pseudonym, which can be used by R to anonymously send a message to S.

5 All messages intended for R will be addressed to $X@S$. These messages will first reach the Forwarding Agent S, and then go through a sequence of FAs to finally reach R, as described below.

Message Forwarding

Once the pseudonym X has been registered with a Forwarding Agent S, messages meant for pseudonym X can be accepted and successfully forwarded to the intended recipient. This takes place as follows.

10 Once the Forwarding Agent S obtains a message intended for X, it looks up X in its internal table. It retrieves two pieces of information from the table; an encrypted version of the address of X (which we will call the “onion address” of X (see Syverson)), as well as the group of FAs to be used for forwarding. The Forwarding Agent S then creates the list of the FAs that the message will pass through; that is, all FAs other than S who will have to “peel the onion” before the address of the intended recipient is revealed. This list contains all the members of the appropriate group except the Forwarding Agent S itself. S affixes this list to the head of the message.

15 The list of FAs needed for X is $S = A_1, A_2, \dots, A_k$. S “peels” one layer of the onion address of X, by applying its private key on the encrypted address, and places this peeled address in the message. S then selects the next forwarding agent at random and forwards the message to it.

25 Every other agent that receives this message performs some verifications to ensure that the protocol has been followed correctly by other agents so far. It then removes one layer of encryption from the onion address received by it. If it is the last agent, it forwards the message to the address thus obtained. Otherwise, it selects an agent at random from the list of agents to be

visited contained in the message header. It then forwards the message to the selected agent after updating certain information, including the list of unvisited agents and the onion address, in the message header.

Encryption

5 All communication between any two Forwarding Agents is encrypted so that an eavesdropper is not able to correlate messages entering and leaving an agent. Such encryption may be dispensed with if it is not considered necessary to guard against global eavesdroppers (for instance, if the FAs are spread far and wide geographically preventing anyone from tracing a message
10 from its point of origin to its terminus).

Each pseudonym shares a secret symmetric key with the FA with which it is registered. When the FA receives a message intended for that pseudonym, it encrypts the message with this symmetric key, after adding a fixed number of random bytes to each encryption block of the message. For
15 example, if the 64-bit DES (Data Encryption Standard) is used, the FA inserts two random bytes before every six bytes of the actual message, and then encrypts the message. The recipient simply discards the first two bytes of every 8-byte block that he receives. The purpose is to defeat attacks in which an adversary sends multiple, identical copies of a message in succession.

20 Detailed Description

At each step, the process outlined above takes place as illustrated in Figure 2. Each agent A, on receiving a message, first verifies that the signed sequence of steps that the message has supposedly gone through is consistent. This is shown in Figure 2 starting at decision block 500, which determines if
25 A is the first agent. If not, execution continues to function block 2000. Otherwise, the appropriate data is retrieved using an internal table lookup in

5

10

15

20

Two Enhancements

The security of this protocol can be enhanced further by two simple modifications to the protocol. Firstly, instead of addressing a message to X as X@P, we could let P itself be a pseudonym, so that the message is actually addressed as X@P@R where R happens to be a real address and X and P are pseudonyms. Once R gets the message, it tunnels it to P using the same strategy outlined earlier and P, in turn, forwards the message to X. Instead of making the sender having to insert the names of the intermediate agents, we make the receiver do the groundwork by registering at multiple spots and setting up the sequence in which the message goes through them. Then the sender does not have to address a message as X@P@R. He just addresses it to X@R. R will send the message to P because X would have given the encrypted address of P while registering with R. When P gets the message, he or she will send it to X because X would have given him or her his or her own encrypted address while registering. This is easily generalized for any number of intermediaries.

For even greater security, the receiver can register at each spot with exactly two pseudonyms, with every two consecutive points sharing one pseudonym and the non-consecutive ones not sharing any pseudonym.

The second enhancement involves a restriction on the number of messages that can be received by any one person within a certain period of time. This is important to prevent a correlation being made on the basis of the number of messages heading for the same destination. The FA with which a person is registered keeps track of the number of messages destined for his or her pseudonyms. Simultaneously, all agents which forward a message to an actual destination also keep track of the number of such messages. If the originating FA is malicious and does not impose this limit, it is likely that one of the endpoints will detect an abnormally high number of messages intended

for a particular address and will bring the system to a halt.

Solution One

The Signature Process

5 In Solution One, we record the path followed by every message through the system, in order to ensure and check that all agents follow the protocol correctly. This path information is stored in the message header. To ensure that it is not modified by dishonest agents, pieces of this information are digitally signed by various agents along the path.

10 This requires a random number N to be chosen as the tag for each message and affixed to it by the Forwarding Agent S , the first agent involved in the forwarding of that message. It is desirable that the Forwarding Agent S should be prevented from choosing the random number N in a non-random manner. This can be done by adopting one of two approaches. The first approach is to use an agreement protocol to select the random number rather than allowing the Forwarding Agent S to choose it himself. A simpler approach would be for each FA to maintain a list of the most recently used random numbers and ensure that the Forwarding Agent S does not try to repeat any of the numbers recently used. This would be sufficient because the protocol simply requires a distinct tag for each message. The statistical distribution of these numbers does not matter.

20 The Forwarding Agent S then selects the next forwarding agent at random. Let the FA thus selected be A_j . The Forwarding Agent S then combines A_j 's name with the tag N and digitally signs the resulting plaintext. Such combining can be in done in many ways, such as by adding or concatenating. It then places the plaintext together with the digital signature created at the head of the list and removes A_j from the list. Finally, the

25

When any other agent A receives this messages, it verifies the integrity of the path information in the header by verifying a succession of signatures. Before forwarding the message to another agent B, A uses the random number N to sign and record the identity of B in the message header. The details of the signing and verification steps are provided in later sections.

10

15

In Solution One, we conceal the identity of the first FA from most of the FAs by splitting the agents into domains. Once the message enters a domain, it visits all the FAs in the domain; when it leaves the domain, there is a marker attached to the head of the message indicating that the particular domain has been visited. All domains except the first are visited in a random sequence. The protocol is followed just as outlined earlier within each domain. Here, the originating agent is known only to the FAs in the initial domain. The only situation in which the final FA in the chain can learn the identity of the originator is if there is a colluding agent in the first domain who broadcasts his

5

10

15

20

Otherwise, the next step in the verification process is the one shown in function block 2040A in Figure 4B, where A finds out the name S of the first agent to receive this message in the current domain. Then, in decision block 2050A, A verifies the signature of S on the first part of the signed sequence in the message header. If this verification succeeds, then in function block 2060A the variable NEXT is initialized to name of the agent that received the message from S (this name is contained in the first part of the signed sequence). Also in function block 2060A, the set SEEN is initialized to contain A and the variable THIS is initialized to S. In decision block 2070A, A checks if there any segments that remain to be verified in the signed sequence. If not, A verifies if the name NEXT and THIS are the names of itself and the agent which forwarded the message to it, respectively, in decision block 2130A, and whether the set SEEN and the list of unvisited agents are disjoint, in decision block 2140A. If both verifications succeed, then the entire verification process is completed and execution continues to function block 2500.

If, on the other hand, more segments are found in decision block 2070A, then in decision block 2080A, A verifies that NEXT is a valid agent name not contained in Z. If so, A adds NEXT to the set SEEN, in function block 2090A, and sets the variable THIS to NEXT, in function block 2100A. In decision block 2110A, A verifies the signature of THIS on the next segment of the signed sequence. If the signature is correct, then in function block 2120A a new agent name NEXT is extracted from that segment. Following this, the execution comes back to decision block 2070A.

If any of the verifications in decision blocks 2050A, 2080A, 2110A, 2130A and 2140A fail, the current message is aborted.

Selecting the Next Agent

In function block 4000 in Figure 2, an agent A which has received a message selects the next agent that will receive this message. Figure 5 illustrates how this selection is performed in Solution One. The process begins at decision block 4010A where A checks if there are any more agents to be visited in the present domain. If not, then A marks the present domain as visited and removes the signed sequence from the message header in function block 4020A. Then in function block 4030A it chooses an unvisited domain at random and makes it the present domain. In function block 4030A, an agent belonging to the current domain is chosen at random from the list of unvisited agents. Following this, the execution continues to function block 4500 in Figure 2.

If, instead, A finds in decision block 4010A that not all the agents in this domain have been visited, then in function block 4050A it chooses at random an unvisited agent belonging to the current domain. It then combines the random number N with the name of this chosen agent and signs the resulting plaintext in function block 4060A. In function block 4070A, this plaintext and signature is added to the signed sequence, following which the execution continues to function block 4500, where the message is forwarded to the chosen agent.

Protocol Correctness

When all the FAs are honest, it is clear that this protocol is double-blind. Let A_i be the last FA in the sequence of FAs involved in forwarding a message meant for a pseudonym X. Let S be the first agent in the sequence. The protocol forwards a message such that it passes through each FA corresponding to the pseudonym exactly once. Thus, once an FA applies its decrypting key on the encrypted address and forwards the message, it does

5

10

15

20

25

Security Against External Attacks

5 A local eavesdropper in this context is an attacker who can observe all
and only communication emanating from a particular individual FA. If the FA
that he or she monitors does not happen to be the last FA in a message
forwarding sequence, then all that the eavesdropper sees is an encrypted
address and a partition of the set of FAs involved in the forwarding into
10 visited and unvisited FAs. Clearly, this information is of absolutely no use to
him. He or she can find neither a pseudonym nor an address. Even if the
monitored FA happens to be the last in a forwarding sequence, the
eavesdropper gathers only an address. He or she may perhaps learn that the
recipient is a client of the system but has no means of ascertaining the
15 pseudonym used by the client. Thus, the system provides absolute privacy
against local eavesdroppers.

A global eavesdropper can potentially attempt to trace the course of a message through the FAs to the final destination. Unlike mixes (see Chaum), we do not provide a direct defense against a global eavesdropper. Tracing on the basis of message content by using a cryptographic system to encrypt communication between any two machines could be prevented. As for traffic analysis, the main defense is to render it difficult by spreading the FAs far and wide across several administrative and geographical domains, making it difficult for a global eavesdropper to exist (see Reiter).

Security Against Malicious FAs

When there are malicious agents in the list of k FAs corresponding to a pseudonym, they still cannot manipulate the list to their advantage by violating the protocol. In particular, they cannot reintroduce the names of FAs already visited so as to enhance the probability that the last FA in the sequence is malicious. Thus the best they can do is to avoid choosing the next agent randomly and instead choose a non-colluder this would eliminate a non-colluding contender for the opportunity to be the last FA in the sequence, thus improving the chances of the malicious agents remaining in the list.

The extent to which single and multiple malicious FAs which can compromise the system security will be further considered.

A single malicious FA

Consider the case when exactly one of the FAs is dishonest. When the FA acts passively, all the FA can gather is the address of some recipients whose messages happen to be forwarded by the FA. But this information is not very useful as the FA cannot associate any pseudonym with the addresses.

But the FA can cause more damage if the FA turns more active. The FA can attempt a known plaintext attack by sending hundreds of copies of the same message addressed to the same pseudonym. One of two scenarios will unfold.

1. The FA does not belong to the list of FAs needed for that particular pseudonym. In this case, there is nothing that the FA can learn. But the recipient will realize that someone is attempting an attack on him and will complain. The pseudonym of the sender can then be blacklisted.

This could prove an effective deterrent to such attacks.

2. The FA is a member of the list corresponding to the particular pseudonym. In this case, about $1/k$ th of the messages will end up with

5

15

15

20

YO999-364

$$P_c = \frac{\left(\frac{c}{n}\right)(c-1)}{(n-1)} = \frac{c(c-1)}{n(n-1)}$$

$$P_{1+} = 1 - \left(\frac{{}^{n-c}C_d}{{}^nC_d} \right)$$
$$P_s \leq P_{1+} \left(\frac{c}{(n-d)} \right) = \left(1 - \left(\frac{{}^{n-c}C_d}{{}^nC_d} \right) \right) \times \frac{c}{(n-d)}$$

YO999-364

*Solution Two*Coordinating Agents

In addition to Forwarding Agents, Solution Two involves another set of agents called Coordinating Agents or CAs. There are r Coordinating Agents labeled with numbers i , $0 \leq i \leq r - 1$. The CAs ensure that the protocol is executed correctly by the FAs and that no malicious FA can tamper with the protocol to break the anonymity that the system promises. The CAs participate in the protocol in the order of their labels. Essentially, they ensure that the number of unvisited FAs contained in the header is decremented correctly every time the message reaches a new FA. It may be noted that though a CA is functionally different from an FA, nothing prevents them from sharing the same physical systems.

The protocol followed by each CA is described in detail in a later section.

15 Signatures

Each message has a header which contains a list of FAs yet to be visited. It also contains the number of such unvisited FAs signed by the CA which has most recently participated in the protocol. As in Solution One, the signature process requires that a tag N for every message. This tag is a random number collectively chosen by all CAs. As in Solution One, the role of the tag is essentially to guard against forgery of signatures.

Verification of Protocol Consistency by an FA

In function block 2000 in Figure 2, a forwarding agent (FA) A verifies if the protocol has been followed consistently so far. In Solution Two, this verification is performed as illustrated in Figure 6. In function block 2010B, A

computes its own sequence number i in the path followed by this message through the set of forwarding agents. This number is computed by subtracting the number of FAs in the list of unvisited FAs from $k + 1$. In decision block 2020B, A checks if i is 1. If i is 1, then A sends CA 0 a request for a tag, in
 5 function block 2080B, and receives the tag N as well as the number $k - 1$, combined with N and signed, in function block 2090B. Following this, the execution continues to function block 2500 in Figure 2.

If the number i is found in decision block 2020B to be different from 1, then in decision block 2030B, A verifies the signature of CA $(i - 2) \bmod r$
 10 on the signed number in the message header. If verification succeeds, then in decision block 2040B A verifies if the signed number is $k + 1 - i$. If the verification succeeds, A sends the numbers $k + 1 - i$ and N and the name of the previous FA to CA $(i - 1) \bmod r$ in function block 2050B. In function block 2060B, A receives a signed number and a signal from that CA. In
 15 decision block 2070B, A verifies if the signal is "OK". If that is the case, verification is complete and execution continues to function block 2500 in Figure 2.

If any of the verifications in decision blocks 2020B, 2030B, 2040B and 2070B fail, A concludes that the protocol has not been executed correctly and
 20 aborts the current message.

Verification of Protocol Consistency by a CA

As described above, verification of protocol consistency requires communication between FAs and CAs. The protocol followed by a CA is illustrated in Figures 7A and 7B.

25 The role of CA 0 is somewhat different than other CAs. This difference is specified in Figure 7A, where in decision block 5100B, a CA (whom we will refer to as C) verifies if it is CA 0. If that is the case, the request must be

5

10

15

5

10

15

20

When all the FAs are honest, it is clear that this protocol is double-blind. Let A_1 be the last FA in the sequence of FAs involved in forwarding a message meant for a pseudonym X . Let S be the first agent in the sequence. The protocol forwards a message such that it passes through each FA corresponding to the pseudonym exactly once. Hence, once an FA applies its decrypting key on the encrypted address and forwards the message, it does not get an opportunity to see the address at any later stage. Therefore, only the last FA in the chain, A_1 , is aware of the actual address. But it may be noted that no individual FA in the chain except for the originator S is aware of the pseudonym to which the message is addressed. Thus, S is aware of the pseudonym X associated with the message, A_1 is aware of the address associated with the message and the intermediate FAs can associate neither with the message. The CAs also do not see either the pseudonym or the address. We thus have a de-linking of pseudonyms from actual addresses. Thus, the goal of receiver anonymity is achieved.

We consider, in order, various types of attacks, ranging from external eavesdroppers, local or global, to attacks by a colluding set of FAs and CAs intended to break the anonymity guaranteed by the system. We will show various defenses against these multifarious attacks and provide probabilistic bounds, where applicable, for the security provided by the system. There are essentially two parameters used for describing this security. The first parameter, which we shall term P_C , is the probability of an identity being compromised as a direct result of multiple colluding FAs being present in the system. The second parameter, P_S , measures the probability that a malicious coterie of colluding FAs and CAs manages to identify the originating FA of a

message whose physical destination is known. This information does not directly endanger the anonymity of any client but may, in the long run, prove useful to attackers in deciding whom to compromise.

We also mention various deterrents to attackers, these being measures which improve the security provided but not amenable to easy quantification.

We thus provide qualitative and quantitative guarantees of security against attacks from within and without.

Security Against External Attacks

The security against external attacks is exactly identical to that of Solution One. Please refer to the corresponding section for Solution One for further details.

Security Against Malicious FAs and CAs

When there are malicious agents in the Forwarding Set (corresponding to a pseudonym) or Coordinating Set, they still cannot manipulate the list to their advantage by violating the protocol. In particular, they cannot reintroduce the names of FAs already visited so as to enhance the probability that the last FA in the sequence is malicious.

Consider the following scenario. Let a malicious FA, F , receive a message in the process of being forwarded. If there were a malicious CA, C , whose ordinal number is t , then F could introduce the names of other malicious agents that have been visited and contact C for the signing procedure. The number of agents introduced would depend on t (CA whose ordinal number is j , has to sign a number which is of the form $k - l * j$, $l = 1, 2, \dots$ where k is the total number of FAs involved in forwarding of a message. This attack fails due to the fact that the CAs communicate with each other at every message forwarding step. Thus, if C were to be contacted by F

5
10

The situation is exactly identical to the corresponding case in Solution One.

15

20

25

If there are c colluders in the collection of n forwarding agents and any agent is equally likely to be a colluder, the probability of an identity being compromised is given by

$$P_c = \frac{\left(\frac{c}{n}\right)(c-1)}{(n-1)} = \frac{c(c-1)}{n(n-1)}$$

We next consider the probability P_s that the identity of the originating FA for a message becomes associated with the physical address of a recipient. For this to happen, the last FA in the sequence has to be malicious. Further, we note that the identity of the originator is known to only three agents—the originator himself, the second FA in the forwarding sequence (since he receives the message from the originator) and CA 0. CA 0 does not see the message itself. So, CA 0 and the last FA cannot identify a correlation on the basis of message content. But the random number associated with the message could be used to do the correlation. In defence, we could periodically reassign the numbers allotted to the CAs so that CA 0 is not statically determined. Then, the probability that the identity of the originator is compromised is simply c/n multiplied by the probability that at least one of the three agents specified is corrupt. This probability P_{1+} , is easily computed as 1 minus the probability that none of these agents are corrupt. Thus,

$$P_{1+} = 1 - \frac{{}^{n-c}C_3}{{}^{n-1}C_3}$$

Therefore,

$$P_s = P_1 + \left(\frac{c}{n}\right) = \left(\frac{c}{n}\right) \left(1 - \frac{{}^{n-c}C_3}{{}^{n-l}C_3}\right)$$

Conclusion

We have presented a new approach to achieving receiver anonymity and have outlined a system that enables a double-blind communication in which neither the sender nor the receiver of the communication is aware of his correspondent's true identity. We have also shown the security of the system against local and global eavesdroppers. We have described various attacks on the system and how they may be repulsed. We have provided qualitative and quantitative measures of the security of the system in the face of malicious agents.

While the invention has been described in terms of preferred embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.